

## **ATTACHMENT 2: Profile and Detailed Achievements of the Group B Recipient of the 2020 C&C Prize**

### **Dr. Michael Stonebraker**

#### **Current Positions:**

Computer Science Professor, Massachusetts Institute of Technology

#### **Personal History (born in 1943):**

- 1965 BSE, Electrical Engineering, Princeton University
- 1967 MS, Electrical Engineering, University of Michigan
- 1971 PhD, Computer, Information and Control Engineering, University of Michigan
- 1971 Assistant Professor, Department of Electrical Engineering and Computer Sciences, University of California, Berkeley
- 1982 Professor, University of California, Berkeley
- 2001 Lecturer, Massachusetts Institute of Technology (MIT), and he later became co-director of the Intel Science and Technology Center for Big Data at MIT's Computer Science and Artificial Intelligence Laboratory
- 2002 Adjunct Professor, Massachusetts Institute of Technology (MIT)

#### **Major Awards:**

- 1988 ACM Software System Award
- 1994 ACM Fellow
- 1992 ACM SIGMOD Edgar F. Codd Innovations Award
- 1994 Association for Computing Machinery (ACM) Fellow
- 1997 National Academy of Engineering (NAE) Member
- 2005 IEEE John von Neumann Medal
- 2014 ACM Turing Award
- 2015 ACM SIGMOD Systems Award

#### **Achievements**

We make use of a variety of different databases on a day-to-day basis. For example, when shopping online, entering your zip code on the shipping form may automatically input the name of the city in which you live. This feature is possible because the form is linked to an address book database. Database systems are also used on websites when you book a flight or look up books on an online library catalog. Companies use computers to manage large amounts of data generated through business activities such as purchasing, production, inventory, sales, employee management, and accounting. This data is managed with a database system to help improve business efficiency. In recent years, efforts have been made to take advantage of this data to resolve societal issues and create new value. For

example, large amounts of data can be analyzed to discover patterns, gain new insights about the data, and implement machine learning so that predictive models can be built to make predictions regarding the future. As the value of data increases, database systems are becoming more and more essential in the efficient management and usage of large amounts of data of all kinds.

Dr. Michael Stonebraker has worked on the research and development of Relational Database Management Systems (RDBMS) and conceived many of the ideas essential to RDBMSs. His work has made a significant impact on many real-world RDBMSs through his involvement in academic prototype publications and startup companies.

Dr. Stonebraker became an assistant professor at the University of California, Berkeley in 1971 and did some early pioneering work with RDBMSs. In the early 1970s, he started the RDBMS research project Ingres after reading a paper published by Edgar F. Codd on the relational model, which is the theoretical basis for relational databases. Along with IBM's System R, Ingres became known as a system that made it possible to implement the relational model in a practical and efficient manner. Many ideas were used in Ingres, including B-trees, which are model tree structures featuring branched diagrams of the data that result in balanced search trees suitable for implementation in databases. There was also the introduction of data replication, which both guarantees data consistency and makes load balancing possible in distributed environments, as well as the introduction of integrity constraints to guarantee data integrity, locking to guarantee the performance of transaction processing, and many other ideas that have been widely adopted in subsequent RDBMSs. By the mid-1970s, Ingres system was functional running on DEC hardware, and by the early 1980s, it was being used commercially. The research results were made available under the derivative work Berkeley Software Distribution (BSD) license agreements, allowing multiple companies to develop products based on BSD. Dr. Stonebraker has also launched various commercial startups.

In the 1980s, in response to the limitations of the relational model, he began developing Postgres, an object-relational database system that would succeed Ingres. Postgres provided an object-relational programming model to support complex data types and allowed users to register new data types as well as program functions that work with these data types. The improvement of both the programmability and the performance of Postgres led to a substantial expansion of the commercial database market. Postgres was distributed under a BSD-style license and became the basis for the free software PostgreSQL available today. A number of companies have also been founded based on PostgreSQL.

In 2001, Dr. Stonebraker moved to the Massachusetts Institute of Technology and began a new series of research projects in response to the changing needs of databases. He developed several database systems including C-Store, a column-oriented database management system, and H-Store, a distributed in-memory database with high throughput. He continues to be actively involved in activities centered on the database field that transcend the boundaries between industry and academia.

Starting in 1988, he worked as the editor of "Readings in Database Systems" (MIT Press), a compilation of important selected papers and works for database experts. This collection of reading is now in its 5<sup>th</sup> edition.

Dr. Stonebraker made tremendous contributions in the early days of database technology by devising numerous RDBMS concepts, demonstrating them, and making them practical, thus laying the foundation for the modern field of database systems. Because Dr. Michael Stonebraker has made so many boundary-transcending achievements in both business and academia, we believe he is a recipient worthy of the C&C Prize.